Periodicity, Repetitions, and Orbits of an Automatic Sequence

Jean-Paul Allouche [a], Narad Rampersad [b], Jeffrey Shallit [c],*

[a]*CNRS, LRI, UMR 8623, Université Paris-Sud, Bâtiment 490, F-91405 Orsay Cedex, France*

[b]*Department of Mathematics and Statistics, University of Winnipeg, 515 Portage Avenue, Winnipeg, MB R3B 2E9, Canada*

[c]*School of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

Abstract

We revisit a technique of S. Lehr on automata and use it to prove old and new results in a simple way. We give a very simple proof of the 1986 theorem of Honkala that it is decidable whether a given k-automatic sequence is ultimately periodic. We prove that it is decidable whether a given k-automatic sequence is overlap-free (or squarefree, or cubefree, etc.) We prove that the lexicographically least sequence in the orbit closure of a k-automatic sequence is k-automatic, and use this last result to show that several related quantities, such as the critical exponent, irrationality measure, and recurrence quotient for Sturmian words with slope α, have automatic continued fraction expansions if α does.

Key words: automatic sequence, squarefree, overlapfree, Thue-Morse sequence, Rudin-Shapiro sequence, decidability, periodicity, orbit, orbit closure, continued fraction

In Honor of Juhani Karhumäki's 60th Birthday

* Corresponding author.
  Email addresses: `allouche@lri.fr` (Jean-Paul Allouche),
`n.rampersad@uwinnipeg.ca` (Narad Rampersad), `shallit@cs.uwaterloo.ca`
(Jeffrey Shallit).

1 Introduction

A sequence $(a_n)_{n \geq 0}$ over a finite alphabet Δ is said to be *k-automatic* for some integer $k \geq 2$ if, roughly speaking, there exists an automaton that, on input n in base k, reaches a state with the output a_n. More formally, a sequence $(a_n)_{n \geq 0}$ over Δ is k-automatic if there exists a deterministic finite automaton with output (DFAO) $M = (Q, \Sigma_k, \Delta, \delta, q_0, \tau)$ where Q is a finite set of states, $\Sigma_k = \{0, 1, 2, \ldots, k-1\}$, $\delta : Q \times \Sigma_k \to Q$ is the transition function, and $\tau : Q \to \Delta$ is the output function, such that if w is any base-k representation of n, possibly with leading zeroes, then $a_n = \tau(\delta(q_0, w^R))$. (Note that $a_0 = \tau(q_0)$.) Here w^R is the reverse of the word w.

This class of sequences, also called k-recognizable in the literature, has been studied extensively (e.g., [9]) and has several different characterizations, the most famous being images (under a coding) of fixed points of k-uniform morphisms.

The archetypal example of a k-automatic sequence is the *Thue-Morse sequence*

$$\mathbf{t} = (t_n)_{n \geq 0} = 0110100110010110 \cdots,$$

where t_n is the sum (modulo 2) of the bits in the base-2 expansion of n [8]. See Figure 1. It can also be viewed as the fixed point of the morphism μ where $0 \to 01$ and $1 \to 10$.
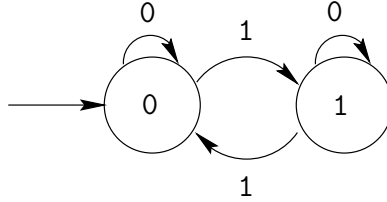

Fig. 1. Automaton generating the Thue-Morse sequence

Given a k-automatic sequence, one might reasonably inquire as to whether the sequence is ultimately periodic. More precisely, we would like to know if the problem

> Given a k-automatic sequence, is it ultimately periodic?

is decidable (i.e., recursively solvable). This problem was solved by Honkala [20], who gave a rather complicated decision procedure.

In this paper, we begin by recalling a technique of Lehr [28] as simplified by Allouche and Shallit [9, pp. 380–382]. In Section 2 we introduce it and use it to reprove the result of Honkala mentioned above.

Another topic of great interest is the pattern-avoiding properties of certain automatic sequences. For example, more than a hundred years ago Thue proved [37,38] that $\mathbf{t}$ contains no overlaps, where an *overlap* is a word of the form $axaxa$, where a is a single letter and x is a word, possibly empty. Examples of overlaps include `alfalfa` in English, `entente` in French, and `ajaja` and `tutut` in Finnish.

Similarly, much attention has been given to avoiding squares. A *square* is a word of the form xx where x is nonempty. Examples of squares include `murmur` in English, `chercher` in French, and `valtavalta` in Finnish. A (finite or infinite) word is *squarefree* if it contains no square factor. As is well-known, if one counts the lengths of the blocks of 1's between consecutive 0's in $\mathbf{t}$, one obtains the squarefree sequence

$$\mathbf{v} = (v_n)_{n \geq 0} = \texttt{210201210120} \cdots .$$

The word $\mathbf{v}$ is generated as the fixed point of the morphism g defined by $2 \to \texttt{210}$, $1 \to \texttt{20}$, and $0 \to \texttt{1}$. Furthermore, $\mathbf{v}$ is generated by the automaton depicted in Figure 2. Here the input is n expressed in base 2, starting with the least significant digit, and the output, given by the symbol labeling the state, is v_n. (Contrast this with the representation given by Berstel [10].)
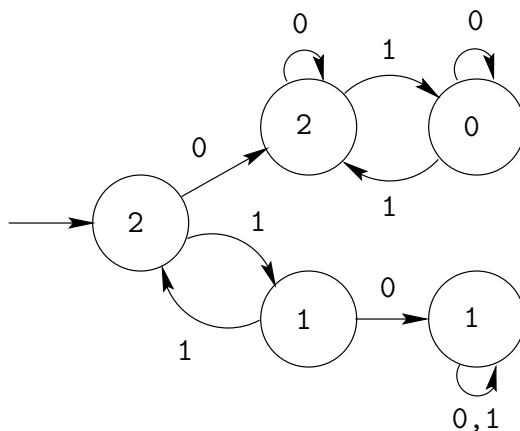

Fig. 2. Automaton generating a squarefree sequence

We can generalize the concept of power to non-integer powers. Let α be a real number > 1. We say that a word z is an α-*power* if it is the shortest prefix of length $\geq \alpha|x|$ of some infinite word $x^\omega = xxx \cdots$, and we say it is an α^+-*power* if it is the shortest prefix of length $> \alpha|x|$ of x^ω. For example, the English word $z = \texttt{abracadabra}$ is both a $3/2$ and a $(3/2)^+$ power, as z is a prefix of length 11 of $(\texttt{abracad})^\omega$, and $10/7 < 3/2 < 11/7$. Using this notation, an overlap is a 2^+ power. We say a (finite or infinite) word z contains an α-power if we can write $z = uvw$ where v is an α-power. We say that a (finite or infinite) word z *avoids α-powers* or is α-*power-free* if it has no factor that is an α-power, and similarly for α^+-powers.

In Section 3 we use Lehr's technique to prove a new result: that it is decidable whether a given k-automatic sequence is squarefree, overlap-free, contains an r-power for r rational, contains an r^+-power, etc.

Let $\mathbf{a} = (a_n)_{n \geq 0}$ be a sequence over a finite alphabet Δ. The *orbit* of $\mathbf{a}$, written $\mathrm{Orb}(\mathbf{a})$, is the set of all its shifts, that is, the set of sequences $\{(a_{n+i})_{n \geq 0} : i \geq 0\}$. The *orbit closure* of $\mathbf{a}$, written $\mathrm{Cl}(\mathrm{Orb}(\mathbf{a}))$ is the closure of $\mathrm{Orb}(\mathbf{a})$ under the usual topology where two sequences are close if they agree on a long prefix. More transparently, a sequence $\mathbf{b} = (b_n)_{n \geq 0}$ is in the orbit closure of $\mathbf{a}$ if and only if every finite prefix of $\mathbf{b}$ is a factor of $\mathbf{a}$ [9, Prop. 10.8.9, p. 327].

An infinite word $\mathbf{a}$ is said to be *recurrent* if every finite factor that occurs in $\mathbf{a}$ occurs infinitely often. It is not hard to see that if $\mathbf{a}$ is recurrent and not periodic, then $\mathrm{Cl}(\mathrm{Orb}(\mathbf{a}))$ is uncountable [9, Thm. 10.8.12, p. 328]. If $\mathbf{a}$ is not recurrent this may not be true; for example, consider the infinite word $\mathbf{c} = \texttt{abaabaaabaaaab} \cdots$. Then $\mathrm{Cl}(\mathrm{Orb}(\mathbf{c}))$ is countable because once a finite factor contains two or more b's, its position in $\mathbf{c}$ is fixed and hence can be extended in at most one way. Thus $\mathrm{Cl}(\mathrm{Orb}(\mathbf{c}))$ equals $\texttt{a}^\omega \cup \texttt{a}^*\texttt{ba}^\omega \cup \mathrm{Orb}(\mathbf{c})$, and hence is countable.

In Section 4 we are interested in elements in the orbit closure of automatic sequences. From the result mentioned above, if $\mathbf{a}$ is recurrent, then "most" of the sequences in $\mathrm{Cl}(\mathrm{Orb}(\mathbf{a}))$ cannot be k-automatic for any k, since the orbit closure is uncountable while the set of k-automatic sequences over Δ is countable. Evidently, this is true even if $\mathbf{a}$ itself is not automatic.

Now suppose that $\mathbf{a}$ is k-automatic, and consider the lexicographically least sequence $\mathbf{b}$ in $\mathrm{Cl}(\mathrm{Orb}(\mathbf{a}))$. We show in Section 4 that $\mathbf{b}$ is also k-automatic, and more generally, any sequence chosen in a periodic way from the factor tree of $\mathbf{a}$ is also k-automatic.

2   Periodicity

Let $\mathbf{a} = (a_n)_{n \geq 0}$ be an infinite sequence. Then $\mathbf{a}$ is *ultimately periodic* if there exist integers $P \geq 1, N \geq 0$ such that $a_i = a_{i+P}$ for all $i \geq N$.

Theorem 1 *Given a DFAO $M = (Q, \Sigma_k, \Delta, \delta, q_0, \tau)$ it is decidable if the k-automatic sequence it generates is ultimately periodic.*

As mentioned before, this result is due to Honkala [20]. We give a new proof.

Proof.   We start with a sketch of the proof. First, we construct an NFA M_1 that on input (P, N) "guesses" I and accepts if $I \geq N$ and $a_I \neq a_{I+P}$. We now convert M_1 to a DFA M_2 using the usual subset construction, and then interchange accepting and non-accepting states, obtaining a DFA M_3 with the property that M_3 accepts (P, N) if and only if $a_I = a_{I+P}$ for all $I \geq N$. Now **a** is ultimately periodic if and only if M_3 accepts some input, which can be checked using the usual depth-first search technique to determine if there is a path from M_3's initial state to a final state.

We now give the proof in detail, addressing concerns such as exactly how P and N are represented, what it means to guess I, how we verify that $I \geq N$, how we compute $I + P$, and what if I is significantly larger than P or N.

When we say that M_1 takes (P, N) as input, what we really mean is that the input alphabet of M_1 is $\Sigma_k \times \Sigma_k$, so that M_1 takes as input the base-k digits of P and N *in parallel*. More precisely, the input is $(p_0, n_0)(p_1, n_1) \cdots (p_j, n_j)$ where $n_j n_{j-1} \cdots n_0$ is a base-k representation of N and $p_j p_{j-1} \cdots p_0$ is a base-k representation of P, either or both padded with leading zeros to ensure that their lengths are the same. This means that (P, N) can be input in infinitely many ways, depending on the number of leading zeros (which are actually trailing zeros since we read the input starting with the least significant digit), and we must ensure that the correct result is returned in each case.

When we say we guess I, what we really mean is that we successively guess the base-k digits of I, starting with the least significant digit.

In order to verify that our guessed I is $\geq N$, we maintain a flag that records how the number represented by the digits of I seen so far stands in relation to the digits of N seen so far: whether it is $<$, $=$, or $>$. The flag is updated as follows, if the next digit of I guessed is i' and the next digit of N is n':

$$
u(<, i', n') = \begin{cases} <, & \text{if } i' \leq n'; \\ >, & \text{if } i' > n'; \end{cases}
$$

$$
u(=, i', n') = \begin{cases} <, & \text{if } i' < n'; \\ =, & \text{if } i' = n'; \\ >, & \text{if } i' > n'; \end{cases} \tag{1}
$$

$$
u(>, i', n') = \begin{cases} <, & \text{if } i' < n'; \\ >, & \text{if } i' \geq n'. \end{cases}
$$

To compute $I + P$, we maintain a "carry" bit, and compute $I + P$ digit-by-digit as we see the digits of P input using the usual pencil-and-paper method.

Finally, since we guess the digits of I in parallel with the digits of the inputs P and N, we have to address the situation where the base-k representation of the appropriate I to guess is longer than the representation of the inputs P and N. If we do not pad P and N with enough 0's, we might return the wrong result. To handle this, we modify the acceptance criterion of the NFA M_1, making a state accepting if an accepting state could be reached by any input of the form $(0,0)^j$, $j \geq 0$.

We now give the construction in more detail. Suppose $M = (Q, \Sigma_k, \Delta, \delta, q_0, \tau)$ is a k-DFAO. We make an NFA $M_1 = (Q', \Sigma_k \times \Sigma_k, \delta', q_0', F')$ as follows.

$Q' = \{<, =, >\} \times \{0, 1\} \times Q \times Q;$
$q_0' = [=, 0, q_0, q_0];$
$F' = \{[b, 0, q, r] \ : \ b \in \{>, =\} \text{ and } \tau(q) \neq \tau(r)\}$

The meaning of a state $[b, c, q, r]$ of Q' is that b is the flag maintaining the relationship between I and N; c is the carry bit in the computation of $I + P$; q is the state in M reached by the bits of I seen so far; and r is the state in M reached by the bits of $I + P$ calculated so far.

We define δ' by $\delta'([b, c, q, r], (n', p')) :=$

$$\{[u(b, i', n'), \lfloor \frac{i' + p' + c}{k} \rfloor, \delta(q, i), \delta(r, (i' + p' + c) \bmod k)] \ : \ 0 \leq i' < k\}.$$

Here u is the update map defined in Eq. (1).

This finishes the construction of the NFA M_1. We now create a new NFA M_1' that is exactly the same as M_1, except that it has a new set of final states $\hat{F}'$ defined by

$$\hat{F}' := \{[b, c, q, r] \ : \quad \text{there exists } j \geq 0 \text{ such that } \delta'([b, c, q, r], (0, 0)^j) \in F'\}.$$

We now convert M_1' to a DFA $M_2 = (Q'', \Sigma_k \times \Sigma_k, \delta'', q_0'', F'')$ using the usual subset construction. We define $M_3 = (Q'', \Sigma_k \times \Sigma_k, \delta'', q_0'', Q'' - F'')$. It is not hard to see that M_3 accepts some input (P, N) with $P \geq 1$ if and only if $\mathbf{a}$ is ultimately periodic. This can be checked by creating a DFA M_4 that accepts $(\Sigma_k^*(\Sigma_k - \{0\})\Sigma_k^*) \times \Sigma_k^*$ and, using the usual direct product construction, creating a DFA M_5 that accepts $L(M_3) \cap L(M_4)$. Then $\mathbf{a}$ is ultimately periodic if and only if M_5 accepts some string, and this can be checked using the usual depth-first search to look for a path connecting the initial state with some final state. ∎

3 Decision problems about repetitions

A morphism $h : \Sigma^* \to \Delta^*$ is said to be k-power-free if whenever w is k-power-free, so is $h(w)$. There is a reasonably large literature about these morphisms, with most investigators concentrating on giving computable characterizations of such morphisms; see, for example, [11,15,21,27,35].

We say a morphism $h : \Sigma^* \to \Sigma^*$ is *prolongable* on a letter a if $h(a) = ax$ for some x such that $h^i(x) \neq \epsilon$ for all $i \geq 0$. In this case there is a unique infinite word with prefixes $h^i(a)$ for all $i \geq 0$, which we write as $h^\omega(a)$. Such a word is called *morphic*. It is also of interest to give computable characterizations of those h for which $h^\omega(a)$ avoids various kind of repetitions. (Note that it is possible for $h^\omega(a)$ to, for example, avoid squares, even if h itself is not squarefree. The morphism g given above in Section 1 provides an example. Here 212 is squarefree, but $g(212)$ is not.)

Berstel [11] showed how to decide if $h^\omega(a)$ is squarefree for three-letter alphabets. Karhumäki [21] showed how to decide if $h^\omega(a)$ is overlap-free for two-letter alphabets. Later, Mignosi and Séébold [31] gave a general algorithm for testing the k-power-freeness of $h^\omega(a)$ for arbitrary non-erasing morphisms h and integers $k \geq 2$. Cassaigne [13] showed how to test if certain kinds of HD0L words avoid arbitrary patterns.

The technique of Section 2 can be modified to create a decision procedure for the existence of many kinds of repetitions in k-automatic sequences. Our approach is both more and less general than previous results in the literature. It is less general because our technique works only for uniform morphisms. It is more general because (a) it works not only for fixed points of uniform morphisms, but also images of those fixed points (under a coding); (b) it works for testing the r-power-freeness and r^+-power-freeness of words, where r is an arbitrary rational number > 1 – a topic relatively unexplored in the literature until now (but see [25,26]); and (c) it works for arbitrary alphabets. We do not know how to make our technique work for r an irrational number.

The following theorem illustrates the technique.

Theorem 2 *The following question is decidable: given a k-automatic sequence $\mathbf{a} = (a_n)_{n \geq 0}$ represented by a DFAO, is $\mathbf{a}$ overlap-free?*

Proof. The proof is very similar to the proof of Theorem 1. The sequence $\mathbf{a} = (a_n)_{n \geq 0}$ contains an overlap if and only if there exist integers $I \geq 0, T \geq 1$ such that $a_{I+J} = a_{I+T+J}$ for all $J, 0 \leq J \leq T$.

Given a DFAO $M = (Q, \Sigma, \Delta, \delta, q_0, \tau)$ for $\mathbf{a}$, we create an NFA M_2 that on input (I, T) accepts if there exists an integer J, $0 \leq J \leq T$, such that $a_{I+J} \neq a_{I+T+J}$. To accomplish this, M_2 guesses the bits of J, verifies that $0 \leq J \leq T$, computes $I + J$ and $I + T + J$ on the fly, and accepts if $a_{I+J} \neq a_{I+T+J}$. As before, we handle the problem that the expansion of $I+T+J$ might be longer than that of I or T by allowing inputs with leading zeroes (actually trailing, since inputs are entered starting with the least significant digit). To do so, we modify the accepting states of M_2 to get a new NFA M_3, by making a state of M_3 accepting if it can be reached in M_2 from an accepting state along a path labeled $(0,0)^j$ for some $j \geq 0$.

We now convert M_3 to a DFA using the subset construction, and change all accepting states to non-accepting and vice versa, obtaining a DFA M_4. Hence M_4 accepts if for all J with $0 \leq J \leq T$ we have $a_{I+J} = a_{I+T+J}$; i.e., there is an overlap of length $2T + 1$ beginning at position I of $\mathbf{a}$. Thus $\mathbf{a}$ contains an overlap if and only if M_4 accepts (I, T) for some integers $I \geq 0$ and $T \geq 1$, which, as before, can be easily checked.

Here are the full details for the construction of $M_2 = (Q', \Sigma_k \times \Sigma_k, \delta', q_0', F')$. The states are 5-tuples of the form $[b, c, d, q, r]$ where b is one of $<, =$, or $>$, expressing the relationship between the guessed J and the input T; c is the carry in the computation of $I+J$; d is the carry in the computation of $I+T+J$; q is the state of M reached on input $I + J$; and r is the state of M reached on input $I + T + J$. The initial state is $q_0' = [=, 0, 0, q_0, q_0]$, and the set of final states is

$$F' = \{[b, 0, 0, q, r] \ : \ b \in \{<, =\} \text{ and } \tau(q) \neq \tau(r)\}.$$

Finally, δ' is defined as follows:

$$\delta'([b, c, d, q, r], (i', t')) = \{ \ [u(b, j', t'), \lfloor \frac{c + i' + j'}{k} \rfloor,$$

$$\lfloor \frac{d + i' + j' + t'}{k} \rfloor, \delta(q, (c+i'+j') \bmod k), \delta(r, (d+i'+j'+t') \bmod k)] \ : \ 0 \leq j' < k\}.$$

∎

Example 3   Using the Grail package [34], version 3.3.4, we verified purely mechanically that the Thue-Morse word $\mathbf{t}$ is overlap-free. We carried out the construction of Theorem 2 by creating an NFA of 72 states (3 possibilities for b, 2 for c, 3 for d (since carries for $d + i' + j' + t'$ could be as much as 2), and 2 possibilities for each of q and r). We added the correct final states, and then converted this to a DFA with 801 states. We then took the complement of this DFA, obtaining a DFA that accepts all pairs (I, T) where there is an overlap of length $2T + 1$ beginning at position I. We then minimized, obtaining a DFA with 2 states that only accepts strings corresponding to $T = 0$. Hence $\mathbf{t}$ is overlap-free.

The same idea can be used to prove each of the following results:

Theorem 4 *Given a DFAO M generating a k-automatic sequence* **a**, *each of the following properties is decidable:*

(a) *Given a rational number* r, *whether* **a** *avoids* r-*powers (resp.,* r^+-*powers);*
(b) *Given a rational number* r, *whether* **a** *contains infinitely many occurrences of* r-*powers (resp.,* r^+-*powers);*
(c) *Given a rational number* r, *whether* **a** *contains infinitely many distinct* r-*powers (resp.,* r^+-*powers);*
(d) *Given a rational number* r, *and a length* l, *whether* **a** *avoids* x^r *(resp.,* r^+-*powers) for* $|x| \geq l$;
(e) *Given a rational number* r, *whether* **a** *avoids* x^r *for all sufficiently long* x;
(f) *Given a length* l, *whether* **a** *avoids palindromes of length* $\geq l$ *(cf. [33]);*
(g) *Whether* **a** *avoids all sufficiently long palindromes;*
(h) *Given a length* l, *whether* **a** *satisfies the property that* x *is a factor of* **a** *of length* $\geq l$, *then its reverse* x^R *is not (cf. [33]);*
(i) *Assuming* **a** *is defined over the alphabet* $\{0, 1, \ldots, j-1\}$, *whether* **a** *avoids all factors of the form* $x\sigma(x)$ *where* $\sigma(a) = (a+1) \bmod j$ *(cf. [29]).*

The proofs for each part are more-or-less trivial variations on the proof of Theorem 2, and we omit them. However, we do make one remark: for parts (a)-(e), we need to replace the condition for the existence of overlaps, namely, "there exist $I \geq 0, T \geq 1$ such that $a_{I+J} = a_{I+T+J}$ for all J, $0 \leq J \leq T$" with the appropriate condition for α-powers, where $\alpha = \frac{p}{q}$ is a rational number. The new condition is "there exist $I \geq 0, T \geq 1$ such that $a_{I+J} = a_{I+T+J}$ for all J, $0 \leq J < (\frac{p}{q} - 1)T$". (In the case of α^+-powers, the inequality becomes $0 \leq J \leq (\frac{p}{q} - 1)T$.) At first sight it might seem difficult to implement this test, for although multiplication can be carried out easily starting with the least significant digit, division is more problematic. To handle this, we simply rewrite the inequality $J < (\frac{p}{q} - 1)T$ as $qJ < (p-q)T$. Now on input T we can guess J digit-by-digit, transduce J into qJ and T into $(p-q)T$, and verify the inequality $qJ < (p-q)T$ on the fly starting with the least significant digit, as before.

4 The orbit closure

We now turn to orbits and the orbit closure of automatic sequences. As motivation, recall that a certain classical dynamical system (i.e., a compact set together with a continuous map of this set) is associated with any sequence, namely the topological closure of the orbit of that sequence under the shift.

For some sequences, the lexicographically least and largest sequences in the orbit closure are known explicitly.

Consider, as an example, the Thue-Morse sequence $\mathbf{t}$. The lexicographically least sequence in the orbit closure of $\mathbf{t}$ is the sequence obtained by iterating the Thue-Morse morphism $\mu : 0 \to 01, 1 \to 10$ on 1, and then dropping the first letter [2,3,5,22]. This gives

$$001011001101001\cdots$$

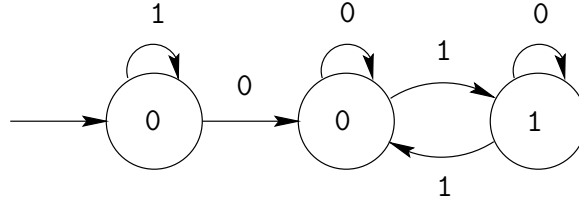and this sequence is clearly 2-automatic, as it is accepted by the DFAO in Figure 3 below.

Fig. 3. Automaton generating the lexicographically least sequence in the orbit closure of the Thue-Morse sequence

Other examples are discussed in Section 6. Recall that the Rudin-Shapiro sequence $\mathbf{u} = (u_n)_{n \geq 0}$ is a 2-automatic sequence defined as follows: u_n is 0 or 1 according to whether the number of (possibly overlapping) occurrences of 11 in the binary expansion of n is even or odd. We observe empirically that the lexicographically least sequence in the orbit closure of the Rudin-Shapiro sequence seems to be the sequence obtained by preceding the Rudin-Shapiro sequence by a 0, but we did not yet prove this.

We now apply the technique of Section 2 to the lexicographically least sequence in the orbit closure of a k-automatic sequence. Our idea is based on the following characterization.

Lemma 5 *Let* $\mathbf{a} = (a_n)_{n \geq 0}$ *be a sequence, and let* $\mathbf{b} = (b_n)_{n \geq 0}$ *be the lexicographically least sequence in the orbit closure of* $\mathbf{a}$. *Then* $b_i = c$ *if and only if there exists* $j \geq 0$ *such that* $a_{j+i} = c$ *and* $a_l a_{l+1} \cdots a_{l+i} \geq a_j a_{j+1} \cdots a_{j+i}$ *for all* $l \geq 0$.

Proof. Suppose $b_i = c$. Then there exists $j \geq 0$ such that $a_j a_{j+1} \cdots a_{j+i} = b_0 b_1 \cdots b_i$, so $a_{j+i} = b_i$. But then $a_l a_{l+1} \cdots a_{l+i} \geq a_j a_{j+1} \cdots a_{j+i}$ for all $l \geq 0$. (Here we use $\geq$ for lexicographic order.)

On the other hand, if $a_l a_{l+1} \cdots a_{l+i} \geq a_j a_{j+1} \cdots a_{j+i}$ for all $l \geq 0$, then $a_j a_{j+1} \cdots a_{j+i}$ must be the prefix of $\mathbf{b}$ of length $i + 1$, and so $b_i = a_{j+i} = c$. $\blacksquare$

10

The advantage to this characterization of b_i is that it does not require explicit knowledge of $b_0, b_1, \ldots, b_{i-1}$.

Theorem 6 *Let* **a** *be* k-*automatic, and let* **b** *be the lexicographically least sequence in the orbit closure of* **a**. *Then* **b** *is* k-*automatic.*

Proof. The idea is to use the condition in Lemma 5. The proof is similar to the proof of Theorem 1, and we outline it below. The fine details about how everything is computed are similar to those of Theorem 1 and we omit them.

The proof consists of several steps. First, suppose we have a k-DFAO M generating **a**. We now create an NFA M_1 that on input (L, J, R) accepts if and only if there exists t, $0 \le t < R$, such that $a_{L+t} \ne a_{J+t}$, or $a_{L+R} \ge a_{J+R}$. The idea is to "guess" t bit-by-bit, verify the inequality $0 \le t < R$, while simultaneously computing the quantities $L + t$, $J + t$, $L + R$, and $J + R$. We accept if $a_{L+t} \ne a_{J+t}$ for some t, $0 \le t < R$, or if $a_{L+R} \ge a_{J+R}$.

From M_1 we create a DFA M_2 that on input (L, J, R) accepts if and only if $a_{L+t} = a_{J+t}$ for all t, $0 \le t < R$ and $a_{L+R} < a_{J+R}$. This is done by converting M_1 to a DFA using the subset construction and changing all accepting states to non-accepting and vice versa. Thus M_2 accepts (L, J, R) if and only if $a_L a_{L+1} \cdots a_{L+R} < a_J a_{J+1} \cdots a_{J+R}$.

Next, from M_2 we create an NFA M_3 that on input (J, R) accepts if and only if there exists an $L \ge 0$ such that $a_L a_{L+1} \cdots a_{L+R} < a_J a_{J+1} \cdots a_{J+R}$. The idea is to "guess" L bit-by-bit and call M_2 on (L, J, R). A priori L could be very big compared to J and R, but our previous trick to handle this works.

Then from M_3 we create a DFA M_4 that on input (J, R) accepts if and only if for all $L \ge 0$ we have $a_L a_{L+1} \cdots a_{L+R} \ge a_J a_{J+1} \cdots a_{J+R}$. This is done by converting M_3 to a DFA using the subset construction, and then changing all accepting states to non-accepting and vice versa.

From M_4 we create an NFA M_5 that on input cI (i.e., the character c concatenated with the base-k expansion of I) accepts if and only if there exists $J \ge 0$ with $a_{J+I} = c$ and $a_L a_{L+1} \cdots a_{L+I} \ge a_J a_{J+1} \cdots a_{J+I}$ for all $L \ge 0$. This is done by recording c in the state, "guessing" J bit-by-bit, computing $J + I$ bit-by-bit and simulating M on $J + I$, and calling M_4 with input (J, I). We then convert M_5 to a DFA M_6 using the subset construction.

Finally, we create a k-DFAO M_7 that on input I simulates M_6 on input cI in parallel for each $c \in \Delta$. Exactly one branch will accept, and the output associated with this branch is c. ∎

5 Continued fraction expansions

The results of the previous section can be generalized to other kinds of orders. Instead of the ordinary lexicographic order, we could consider an order that depends on the index of the string being compared. One way to do this is to consider a sequence of permutations $(\psi_i)_{i \geq 0}$, where each $\psi_i :$ $\Delta \to \Delta$, and when comparing $a_0 a_1 \cdots a_{i-1}$ to $b_0 b_1 \cdots b_{i-1}$, we instead compare $\psi_0(a_0) \cdots \psi_{i-1}(a_{i-1})$ to $\psi_0(b_0) \cdots \psi_{i-1}(b_{i-1})$ (using the ordinary lexicographic order). An example of this kind of ordering comes from continued fractions, where $[a_0, a_1, a_2, \ldots] < [b_0, b_1, b_2, \ldots]$ if and only if $a_0 < b_0$, or $a_0 = b_0$ and $a_1 > b_1$, or $a_0 = b_0$, $a_1 = b_1$, and $a_2 < b_2$, etc. This corresponds to inverting the order of the elements being compared on the odd indexes. Provided the sequence $(\psi_i)_{i \geq 0}$ is k-automatic, the result of Theorem 6 still holds.

Corollary 7 *Let $(\psi_i)_{i \geq 0}$ be a k-automatic sequence of permutations, and let $(a_i)_{i \geq 0}$ be a k-automatic sequence. Then the lexicographically least sequence in the orbit closure, as modified by the permutations (ψ_i), is k-automatic.*

Proof. In the construction of Theorem 6, when we compare a_{L+R} to a_{J+R}, we instead compare $\psi_R(a_{L+R})$ to $\psi_R(a_{J+R})$. Since $(\psi_i)_{i \geq 0}$ is k-automatic, there is no problem computing ψ_R on input R. ∎

From now on, when we talk about a continued fraction expansion $[a_0, a_1, \ldots]$ being k-automatic, we mean the continued fraction has bounded partial quotients and the underlying sequence of partial quotients $(a_i)_{i \geq 0}$ is k-automatic.

Let $T(x)$ be the usual transformation on continued fractions defined by $T(x) = \frac{1}{x - \lfloor x \rfloor}$, so that $T([a_0, a_1, a_2, \ldots]) = [a_1, a_2, \ldots]$. Thus we have

Theorem 8 *Let x be an irrational real number with a k-automatic continued fraction expansion $[a_0, a_1, \ldots]$. Then the continued fraction expansions of both $\liminf_{n \to \infty} T^n(x)$ and $\limsup_{n \to \infty} T^n(x)$ are k-automatic.*

Proof. Use Corollary 7, where the permutations invert the order of the letters on every other index. ∎

In addition to the orbit closure of a sequence, we can study a related structure, which we call the *reverse orbit closure*. We say that a sequence $\mathbf{b} = (b_n)_{n \geq 0}$ is in the reverse orbit closure of $\mathbf{a} = (a_n)_{n \geq 0}$ if every finite prefix of $\mathbf{b}$ is a prefix of some word of the form $a_r a_{r-1} a_{r-2} \cdots a_1 a_0$.

Theorem 9 *If* $\mathbf{a} = (a_n)_{n \geq 0}$ *is k-automatic, then so is the lexicographically least sequence in the reverse orbit closure.*

Proof. Let $b = (b_n)_{n \geq 0}$ be the lexicographically least sequence in the reverse orbit closure of $\mathbf{a} = (a_n)_{n \geq 0}$. We use the following characterization of $\mathbf{b}$: $b_i = c$ if and only if there exists $r \geq i$ such that $a_{r-i} = c$ and $a_s a_{s-1} \cdots a_{s-i} \geq a_r a_{r-1} \cdots a_{r-i}$ for all $s \geq i$.

We can now implement this test in exactly the same way that we implemented the test in the proof of Theorem 6. ∎

We can also combine the reverse orbit closure with a permutation that inverts the order of the letters on every other index.

Theorem 10 *Let α be an irrational real number with a k-automatic continued fraction expansion $[a_0, a_1, a_2, \ldots]$. Let p_n/q_n be the n'th convergent to the continued fraction to α. Let $\beta = \liminf_{n \to \infty} p_n/p_{n-1}$ and $\gamma = \liminf_{n \to \infty} q_n/q_{n-1}$, $\delta = \limsup_{n \to \infty} p_n/p_{n-1}$, $\zeta = \limsup_{n \to \infty} q_n/q_{n-1}$. Then the continued fraction expansion of each of $\beta, \gamma, \delta, \zeta$ is k-automatic.*

Proof. We prove the result for β, the others being similar. By a famous result of Galois [18] we have

$$\frac{p_n}{p_{n-1}} = [a_n, a_{n-1}, \ldots, a_0].$$

Now β corresponds to the lexicographically least sequence in the reverse orbit closure of $(a_i)_{i \geq 0}$, except that the ordering is slightly different from the usual ordering, where the ordering is as usual on the even indexed terms and opposite on the odd-indexed terms. As in Corollary 7, we can handle this in the same way. ∎

Example. Let us consider an example. As is well-known [36,39], for integers $k \geq 3$ the real number

$$\alpha_k = \sum_{i \geq 0} k^{-2^i} = [0, k-1, k+2, k, k, k-2, k, k+2, k, k-2, k+2, k, k-1, \ldots]$$

has a 2-automatic continued fraction expansion, generated by the automaton given in Figure 4 (again, the automaton expects the least significant digit first).
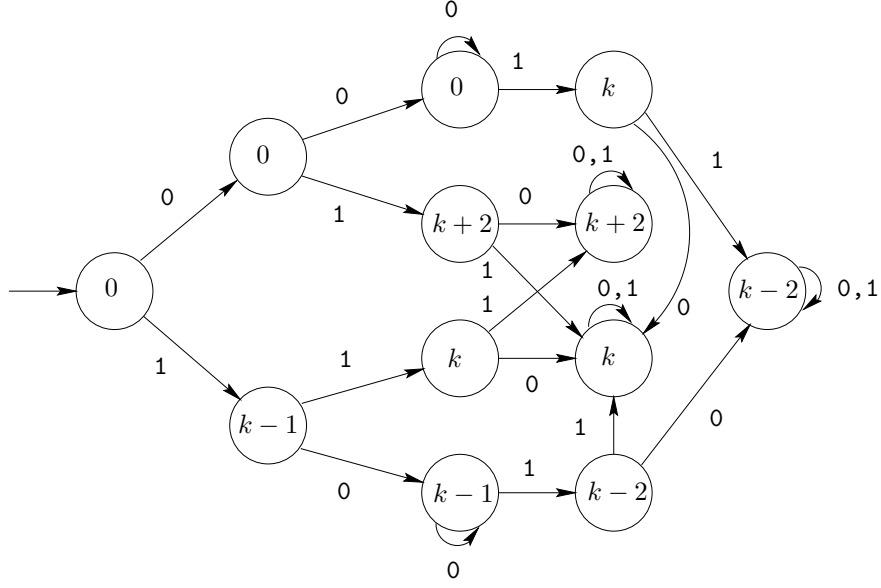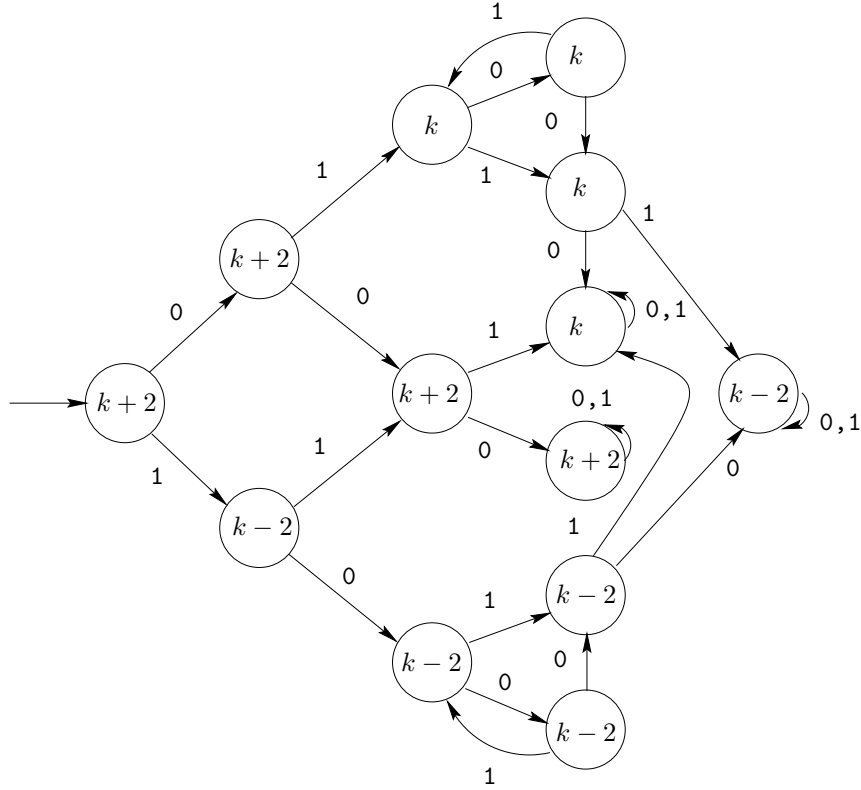
Fig. 4. Automaton generating the continued fraction for α_k

Fig. 5. Automaton generating the continued fraction for ζ_k

Then $\zeta_k = \limsup_{n \geq 0} q_n/q_{n-1} = [k+2, k-2, k, k+2, k, k-2, k, k, \ldots]$ is 2-automatic. ∎

Let α be an irrational number with partial quotients p_n/q_n. The quantity $\zeta = \limsup_{n\geq 0} q_n/q_{n-1}$ figures in a number of recent papers in combinatorics on words. For example, $2 + \zeta$ is the value of the recurrence quotient of a Sturmian word with slope α [14,1]. Hence this recurrence quotient has a k-automatic continued fraction if α does.

The number ζ also appears (actually, $\zeta + 1$) as the irrationality measure of numbers of the form $(b - 1) \sum_{n\geq 1} b^{-\lfloor n\alpha \rfloor}$ [1].

Finally, ζ also appears in a formula giving the critical exponent (aka "index") of Sturmian words, as found by Damanik and Lenz [16, Thm. 1, p. 24] and Cao and Wen [12, Thm. 9, p. 380]. This exponent is essentially

$$\zeta' := 2 + \limsup_{n\geq 1} \frac{q_n - 2}{q_{n-1}}.$$

If the lim sup is actually attained for a particular n, then the critical exponent is rational. Otherwise it clearly coincides with $2+\zeta$, and its continued fraction expansion is k-automatic if that of α is.

6 Applications

Our results about the lexicographically least and largest sequences in the orbit closure of a sequence can be illustrated by and applied to two families of binary sequences: the sequences in the set Γ described below and the Sturmian sequences.

6.1 Sequences in the set Γ

Theorem 6 can be applied to shed some light on the automatic sequences that belong to two sets of binary sequences: the set Γ occurring in the study of iterations of continuous unimodal maps of the interval (see [3,2]) and the set Γ_{strict} occurring in the study of unique β-expansions of the number 1 [17,22,4], where

$$\Gamma \quad := \{A \in \{0,1\}^{\omega} : \ \forall k \geq 0, \ \overline{A} \leq \sigma^k A \leq A\}$$

$$\Gamma_{\text{strict}} := \{A \in \{0,1\}^{\omega} : \ \forall k \geq 1, \ \overline{A} < \sigma^k A < A\}.$$

Here $A = (a_n)_{n\geq 0}$, and σ is the shift on sequences defined by $\sigma A := (a_{n+1})_{n\geq 0}$. The bar operation replaces 0's by 1's and 1's by 0's, i.e., $\overline{A} := (1 - a_n)_{n\geq 0}$. Note that these two sets differ only by a set of (purely) periodic sequences.

15

Also note that the set Γ above differs slightly from the set Γ in [2], in that the set Γ above contains the extra sequence $(10)^\omega$.

The shifted Thue-Morse sequence is an element of Γ, as are more general automatic sequences (e.g., analogues of the Thue-Morse sequence including the q-mirror sequences introduced in [3,2]; see [23,24,40,32,6]).

Now for any binary sequence A belonging to Γ, define, as in [3,2],

$$\Gamma_A := \{B \in \{0,1\}^\omega : \ \forall k \geq 0, \ \overline{A} \leq \sigma^k B \leq A\}.$$

Of course, the sequence A belongs to Γ_A. Furthermore 1^ω belongs to Γ, and any binary sequence B belongs to Γ_{1^ω}. Thus, given B, it is interesting to look for the lexicographically least sequence A such that B belongs to Γ_A. The answer is easy (see [2, pp. 37–38]): the least sequence A in Γ such that B belongs to Γ_A is

$$\Theta(B) := \sup(\{\sigma^k B : \ k \geq 0\} \cup \{\sigma^\ell \overline{B} : \ \ell \geq 0\}).$$

In particular for any sequence B, all sequences $\sigma^k B$ and $\sigma^\ell \overline{B}$ belong to $\Gamma_{\Theta(B)}$, and $\Theta(B)$ is the largest such sequence.

Theorem 6 above shows that if B is automatic, then so is $\Theta(B)$. This remark is a small step in the study of *all* automatic sequences belonging to Γ. Note that Γ is not countable (see e.g., [2, Prop. 3, p. 35]), so that Γ also contains sequences that are *not* automatic. Even more, Γ contains sequences whose subword complexity is not $O(n)$: it suffices to take the sequence $\Theta(B)$, where B is, as in [19], a binary minimal sequence with positive topological entropy, hence with subword complexity not of the form $O(n)$.

6.2 Sturmian sequences

We suppose that the reader is familiar with the notion of Sturmian sequence (see, e.g., [30, Chapter 2]). A result on characteristic Sturmian sequences and Sturmian sequences that was proved or partly proved several times (see the survey [7]) states that

Theorem 11

(a) *A nonperiodic sequence A is characteristic Sturmian if and only if for any $k \geq 0$ the following inequalities hold*

$$0\,A \leq \sigma^k A \leq 1\,A.$$

(b) A nonperiodic binary sequence A is Sturmian if and only if there exists a binary sequence B such that for any $k \geq 0$ the following inequalities hold

$$\texttt{0}\, B \leq \sigma^k A \leq \texttt{1}\, B.$$

Furthermore such a B is unique, and is the characteristic Sturmian sequence having the same slope as A.

Theorem 11 easily implies the following corollary.

Corollary 12 *The lexicographically least (resp. largest) sequence in the orbit closure of a Sturmian sequence A is the sequence $\texttt{0}B$ (resp. $\texttt{1}B$) where B is the characteristic sequence with the same slope as A.*

Proof. It is not difficult to see that the inequalities above are optimal in the sense that, e.g., for a characteristic sequence A, we have $\texttt{0}A = \inf\{\sigma^k A : k \geq 0\}$ and similarly for the other three inequalities in Theorem 11 above. ∎

7 Acknowledgments

References

[1] B. Adamczewski and J.-P. Allouche. Reversals and palindromes in continued fractions. *Theoret. Comput. Sci.* **380** (2007), 220–237.

[2] J.-P. Allouche. *Théorie des Nombres et Automates.* Thèse d'État, 1983. Text available electronically at
`http://tel.archives-ouvertes.fr/tel-00343206/fr/` .

[3] J.-P. Allouche and M. Cosnard. Itérations de fonctions unimodales et suites engendrées par automates. *C. R. Acad. Sci. Paris* **296** (1983), 159–162.

[4] J.-P. Allouche and M. Cosnard. Non-integer bases, iteration of continuous real maps, and an arithmetic self-similar set. *Acta Math. Hung.* **91** (2001), 325–332.

[5] J.-P. Allouche, J. Currie, and J. Shallit. Extremal infinite overlap-free binary words. *Electronic J. Combinatorics* **5**(1) (1998), #R27 (electronic),
`http://www.combinatorics.org/Volume_5/Abstracts/v5i1r27.html`

[6] J.-P. Allouche and C. Frougny. Univoque numbers and an avatar of Thue-Morse. Preprint, http://arxiv.org/abs/0712.0102, 2007. *Acta. Arith.*, to appear.

[7] J.-P. Allouche and A. Glen. Extremal properties of (epi)sturmian sequences and distribution modulo 1. In preparation, 2008.

[8] J.-P. Allouche and J. O. Shallit. The ubiquitous Prouhet-Thue-Morse sequence. In C. Ding, T. Helleseth, and H. Niederreiter, editors, *Sequences and Their Applications, Proceedings of SETA '98*, pp. 1–16. Springer-Verlag, 1999.

[9] J.-P. Allouche and J. Shallit. *Automatic Sequences: Theory, Applications, Generalizations.* Cambridge University Press, 2003.

[10] J. Berstel. Sur la construction de mots sans carré. *Sém. Théor. Nombres Bordeaux*, 1978–79, Exposé 18, 18.01–18.15.

[11] J. Berstel. Sur les mots sans carré définis par un morphisme. in H. A. Maurer, ed., Proc. 6th Int'l Conf. on Automata, Languages, and Programming, ICALP '79, Lect. Notes in Comp. Sci., Vol. 71, Springer-Verlag, 1979, pp. 16–25.

[12] W.-T. Cao and Z.-Y. Wen. Some properties of the factors of Sturmian sequences. *Theoret. Comput. Sci.* **304** (2003), 365–385.

[13] J. Cassaigne. An algorithm to test if a given circular HD0L-language avoids a pattern. Information processing '94, Vol. I, (Hamburg, 1994), *IFIP Trans. A Comput. Sci. Tech.* **A-51**, North-Holland, Amsterdam, 1994, pp. 459–464. DOI 10.1.1.9.2125.

[14] J. Cassaigne. Limit values of the recurrence quotient of Sturmian sequences. *Theoret. Comput. Sci.* **218** (1999), 3–12.

[15] M. Crochemore. Sharp characterizations of squarefree morphisms. *Theoret. Comput. Sci.* **18** (1982), 221–226.

[16] D. Damanik and D. Lenz. The index of Sturmian sequences. *European J. Combinatorics* **23** (2002), 23–29.

[17] P. Erdős, I. Joó, and V. Komornik. Characterization of the unique expansions $1 = \sum_{i=1}^{\infty} q^{-n_i}$, and related problems. *Bull. Soc. Math. France* **118** (1990), 377–390.

[18] E. Galois. Démonstration d'un théorème sur les fractions continues périodiques. *Ann. Math. Pures et Appl.* **19** (1828–9), 294–301.

[19] C. Grillenberger. Construction of strictly ergodic systems. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **25** (1973), 323–334.

[20] J. Honkala. A decision method for the recognizability of sets defined by number systems. *RAIRO Inform. Théor. App.* **20** (1986), 395–403.

[21] J. Karhumäki. On cube-free ω-words generated by binary morphisms. *Disc. Appl. Math.* **5** (1983), 279–297.

[22] V. Komornik and P. Loreti. Unique developments in non-integer bases. *Amer. Math. Monthly* **105** (1998), 636–639.

[23] V. Komornik and P. Loreti. Subexpansions, superexpansions and uniqueness properties in non-integer bases. *Period. Math. Hungar.* **44** (2002), 197–218.

[24] V. Komornik and P. Loreti. On the topological structure of univoque sets. *J. Number Theory* **122** (2007), 157–183.

[25] D. Krieger. On critical exponents in fixed points of k-uniform binary morphisms. *RAIRO-Theor. Inf. Appl.* **43** (2009), 41–68. Corrigenda available at `http://www.wisdom.weizmann.ac.il/~daliak/papers/UniformBinaryThesis.pdf`.

[26] D. Krieger. On critical exponents in fixed points of non-erasing morphisms. *Theoret. Comput. Sci.* **376** (2007), 70–88.

[27] M. Leconte. A characterization of power-free morphisms. *Theoret. Comput. Sci.* **38** (1985), 117–122.

[28] S. Lehr. Sums and rational multiples of q–automatic sequences are q–automatic. *Theoret. Comput. Sci.* **108** (1993), 385–391.

[29] J. Loftus, J. Shallit, and M.-w. Wang. New problems of pattern avoidance. In *Developments in Language Theory (DLT '99)*, pp. 185–199. World Scientific, 2000.

[30] M. Lothaire. *Algebraic Combinatorics on Words*. Cambridge University Press, 2002.

[31] F. Mignosi and P. Séébold. If a D0L language is k-power free then it is circular. In A. Lingas, R. Karlsson, and S. Carlsson, eds., Proc. 20th Int'l Conf. on Automata, Languages, and Programming, ICALP '93, Lect. Notes. in Comp. Sci., Vol. 700, Springer-Verlag, 1993, pp. 507–518.

[32] M. Niu and Z.-x. Wen. A property of m-tuplings Morse sequence. *Wuhan Univ. J. Nat. Sci.* **11** (2006), 473–476.

[33] N. Rampersad and J. Shallit. Words avoiding reversed subwords. *J. Combin. Math. Combin. Comput.* **54** (2005), 157–164.

[34] D. Raymond and D. Wood. Grail: a C++ library for automata and expressions. *J. Symbolic Comput.* **17** (1994), 341–350.

[35] G. Richomme and F. Wlazinski. Existence of finite test-sets for k-power-freeness of uniform morphisms. *Disc. Appl. Math.* **155** (2007), 2001–2016.

[36] J. O. Shallit. Simple continued fractions for some irrational numbers. *J. Number Theory* **11** (1979), 209–217.

[37] A. Thue. Über unendliche Zeichenreihen. *Norske vid. Selsk. Skr. Mat. Nat. Kl.* **7** (1906), 1–22. Reprinted in *Selected Mathematical Papers of Axel Thue*, T. Nagell, editor, Universitetsforlaget, Oslo, 1977, pp. 139–158.

[38] A. Thue. Über die gegenseitige Lage gleicher Teile gewisser Zeichenreihen. *Norske vid. Selsk. Skr. Mat. Nat. Kl.* **1** (1912), 1–67. Reprinted in *Selected Mathematical Papers of Axel Thue*, T. Nagell, editor, Universitetsforlaget, Oslo, 1977, pp. 413–478.

[39] A. J. van der Poorten and J. O. Shallit. Folded continued fractions. *J. Number Theory* **40** (1992), 237–250.

[40] M. de Vries and V. Komornik. Unique expansions of real numbers. Preprint, http://arxiv.org/abs/math/0609708v3, 2007. *Adv. in Math.*, to appear.